

Wavefunction preparation and resampling using a quantum computer

Alexei Kitaev and William A. Webb
California Institute of Technology, Pasadena, CA, 91125-3600, USA

We present an algorithm that prepares multidimensional Gaussian wavefunctions on qubit arrays and an application of such wavefunctions to multidimensional resampling, a technique useful in quantum digital simulation.

INTRODUCTION

Shor's algorithms for factoring and discrete logarithms [1], as well as Grover's search algorithm [2] promise substantial quantum speed-up for certain important problems. Another exciting application area for quantum computers is simulation of physical systems, including molecules, strongly correlated spin systems, nuclei, and QCD. The majority of these problems admit a natural description using continuous variables, such as particle coordinates or field modes. ("Mode" may refer to the field strength at a grid point or to a Fourier component, depending on which representation is chosen.) However, we want to perform the simulation using a *digital* quantum computer, i.e., one composed of qubits. In this paper, we address some general issues pertaining to the discretization of continuous variables and round-off errors in the quantum case. It is worth noting that continuous variables play an important role in algebraic number theory (for example, consider the use of logarithmic coordinates in the proof of Dirichlet's unit theorem [3].) Thus, we anticipate that discretization techniques might also be helpful in constructing new quantum algorithms for number-theoretic problems.

To see why the discretization is not entirely trivial, consider a quantum particle characterized by a single coordinate x . To represent it in a quantum computer, we approximate x by $\tilde{x} = j\delta$, where j is a (binary-encoded) integer and δ is some small constant (the discretization grid spacing). To simplify the notation, we may rescale x so that $\delta = 1$; thus, we approximate real numbers by integers. Now, suppose we want to change this standard grid to a different one, which is equivalent to stretching or squeezing the wavefunction and putting it back on the same grid. We call this transformation *resampling*. If x is mapped to $f(x)$, a continuous wavefunction would transform as follows:

$$\psi(x) \mapsto \xi(y) = (f'(x))^{-1/2} \psi(y), \quad \text{where } x = f^{-1}(y). \quad (1)$$

In the simplest case, $f(x) = \frac{x}{a}$, so that $x \mapsto \frac{x}{a}$ and $\psi(x) \mapsto \sqrt{a} \psi(ay)$. Naively, this appears to be implementable as a permutation of basis vectors. However, the discrete version of f ,

$$\tilde{f}(j) = \lfloor f(j) \rfloor = \lfloor \frac{j}{a} \rfloor \quad (2)$$

is not a permutation! For $a < 1$ the transformation is a stretching; it maps integers to a subset of integers, leaving some gaps. For $a > 1$ it is a squeezing, therefore two distinct integers may be mapped to the same integer.

Fortunately, transformation (2) is to be applied to a special class of wavefunction, namely, ones that vary at distances much larger than the grid spacing. In this case, a few lowest bits of x are almost in the uniform superposition and thus can be easily removed or added again while maintaining quantum coherence. But adding or removing k bits corresponds to stretching (resp. squeezing) by factor 2^k . It is easy to generalize this trick to any scaling factor a , or to an arbitrary function f .

Multidimensional resampling is somewhat more challenging though. Our solution for this case uses multidimensional Gaussian states. Of course, the preparation of such states begins with the one-dimensional case. The transition to higher dimensions requires some care. One may be tempted to say that a general Gaussian is just a product of several one-dimensional Gaussians. However, the implementation of this approach involves the diagonalization of the covariance matrix and a linear change of variables, which in turn requires resampling. To avoid this problem, we use a special linear transformation, namely shearing, and construct its discrete version explicitly. In the rest of the paper, we simply expand this outline.

PREPARING A GAUSSIAN IN ONE DIMENSION

The states we wish to prepare are discrete approximations to continuous Gaussian functions. In this section we examine these functions and the approximate forms of interest. In one dimension, the Gaussian is just the normal

distribution familiar from statistics: $G(x, \sigma, \mu) = e^{-\frac{(x-\mu)^2}{2\sigma^2}}$ (up to a normalization factor). Our algorithm is similar to the uniform superposition preparation technique [4] and its generalization to efficiently integrable distributions [5]. However, the latter procedures begin with dissecting the distribution at $x_0 = 2^{N-1}$ and setting the highest bit of x according to $\int_0^{x_0} |\psi(x)|^2 dx$. In comparison, we first set the *lowest* bit and proceed with generating a discrete Gaussian on either even integers or odd integers. This method is slightly more efficient since it does not involve the Gauss error function, but rather, the Jacobi theta function, which is easier to calculate.

If we treat the Gaussian wavefunction as continuous, it has form

$$\psi_{\sigma,\mu}(x) = \frac{1}{\sqrt{\sigma\sqrt{\pi}}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (3)$$

so that

$$\psi_{\sigma,\mu}^2(x) = \frac{1}{\sigma\sqrt{\pi}} e^{-\frac{(x-\mu)^2}{\sigma^2}} \quad (4)$$

and

$$\int_{-\infty}^{\infty} \psi_{\sigma,\mu}^2(x) dx = 1 \quad (5)$$

If we consider it to be a function of an integer value, we have

$$\tilde{\psi}_{\sigma,\mu}(i) = \frac{1}{\sqrt{f(\sigma,\mu)}} e^{-\frac{(i-\mu)^2}{2\sigma^2}} \quad (6)$$

where the function f is related to the third Jacobi theta function (θ_3) and is defined as

$$f(\sigma,\mu) = \sum_{n=-\infty}^{\infty} e^{-\frac{(n-\mu)^2}{\sigma^2}} \quad (7)$$

so that

$$\sum_{i=-\infty}^{\infty} \tilde{\psi}_{\sigma,\mu}^2(i) = 1. \quad (8)$$

If $\sigma \gg 1$, then $f(\sigma,\mu)$ is approximately equal to $\int_{-\infty}^{\infty} e^{-\frac{(x-\mu)^2}{\sigma^2}} dx = \sigma\sqrt{\pi}$. We now define the function

$$\xi_{\sigma,\mu,k}^2(i) = \sum_{j=-\infty}^{\infty} \tilde{\psi}^2(i + j \cdot 2^k) \quad (9)$$

Note that this function is periodic with period 2^k . This allows us to define the quantum state

$$|\xi_{\sigma,\mu,N}\rangle = \sum_{i=0}^{2^N-1} \xi_{\sigma,\mu,N}(i) |i\rangle \quad (10)$$

For $\mu, 2^N - \mu \gg \sigma \gg 1$, this state is very close to a Gaussian with mean μ and standard deviation σ . A recursive description of the state,

$$|\xi_{\sigma,\mu,N}\rangle = \left| \xi_{\frac{\sigma}{2}, \frac{\mu}{2}, N-1} \right\rangle \otimes \cos \alpha |0\rangle + \left| \xi_{\frac{\sigma}{2}, \frac{\mu-1}{2}, N-1} \right\rangle \otimes \sin \alpha |1\rangle, \quad (11)$$

where

$$\alpha = \cos^{-1} \sqrt{f\left(\frac{\sigma}{2}, \frac{\mu}{2}\right) / f(\sigma, \mu)}, \quad (12)$$

leads almost immediately to an algorithm for preparing it. Note that the sum represented by $f(\sigma,\mu)$ converges, so we can approximate it with a finite number of terms and keep some number of digits of the answer. However, the

convergence is poor when σ is large. When σ is small (which is the case at a late stage of our recursive algorithm), a few terms of the sum will dominate, and only these need be calculated and stored. For large σ we use $f(\sigma, \mu) = \sigma\sqrt{\pi} \exp\left(-\frac{\mu^2}{\sigma^2}\right) f\left(\frac{1}{\pi\sigma}, -i\frac{\mu}{\sigma^2\pi}\right)$ to avoid having to sum an exponentially large number of terms.

The algorithm recursively calls the following subroutine, using the input data $(\sigma_0, \mu_0) = (\sigma, \mu)$ and some number of supplementary qubits. Note that the N qubits are indexed 0 to $N-1$, where the 0-th qubit corresponds to the rightmost qubit in a binary representation.

1. (a) Calculate and store α (see eqn. (12)) and $\alpha/2\pi$ in a k -bit binary representation on supplementary qubits.
 (b) Apply the rotation operator $R(\alpha)$ to the 0-th qubit

$$R(\alpha) = \begin{bmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{bmatrix} \quad (13)$$

using the bits of $\alpha/2\pi$ as described below.

- (c) Uncompute α and $\alpha/2\pi$.
2. Compute (σ_1, μ_1) , where $\sigma_1 = \sigma_0/2$ and $\mu_1 = \mu_0/2$ if the previously rotated qubit is $|0\rangle$ and $\mu_1 = (\mu_0 - 1)/2$ if it is $|1\rangle$. Note that all pairs (σ_i, μ_i) are stored in supplementary memory until a cleanup phase, see below.
3. On the remaining $N-1$ qubits prepare state $|\xi_{\sigma_1, \mu_1, N-1}\rangle$.

Note that for the last (($N-1$))th qubit, we proceed only through step 1 (c), as after this qubit is rotated we do not need another pair of parameters. After this last qubit is rotated, right before termination, a last series of steps is performed in which all stored pairs (σ_i, μ_i) are cleared. To do this, the machine erases $(\sigma_{N-1}, \mu_{N-1})$ by recalculating it using $(\sigma_{N-2}, \mu_{N-2})$ and the last output qubit, performing a second controlled XOR to the region where $(\sigma_{N-1}, \mu_{N-1})$ is stored, thus clearing $(\sigma_{N-1}, \mu_{N-1})$ from memory. The process is repeated on $(\sigma_{N-2}, \mu_{N-2})$, $(\sigma_{N-3}, \mu_{N-3})$ and so on until the supplementary memory is cleared.

Since quantum recursion can be counterintuitive and difficult to describe verbally, this algorithm is perhaps easier to understand when presented in circuit form. The figure shows the gist of the algorithm as well as some of the workings of the supplementary qubits. Note that the double line across the top denotes the classical part of the computer, which in this scheme functions only to transfer the specifications of the Gaussian to the quantum part at the end of the process to clear memory. The single lines with forward slashes (/) overlying their left ends represent multiple qubits in the supplementary memory. Note that each such line represents the same array at different stages of the computation. We deviate slightly from conventional notation by introducing the downward arrow conditional, which represents a generalized form of controlled operation.

The downward arrow terminating on the XOR symbol on the supplementary array directly above the i th qubit denotes that the previously used pair $(\sigma_{i-1}, \mu_{i-1})$ and the $(i-1)$ th qubit (which was just rotated) are used to control the calculation of (σ_i, μ_i) and α_i . Note that classically, α_i is a rotation angle equal to $\sqrt{f(\sigma_{i-1}/2, \mu_{i-1}/2)/f(\sigma_{i-1}, \mu_{i-1})}$ if the $(i-1)$ -th qubit is $|0\rangle$ and $\sqrt{f(\sigma_{i-1}/2, (\mu_{i-1} - 1)/2)/f(\sigma_{i-1}, \mu_{i-1})}$ if the $(i-1)$ -th qubit is $|1\rangle$. However, since the previous qubit is always in a superposition state, α_i is always entangled with the previous qubits.

The downward arrow terminating on the i -th rotation operator box indicates that the rotation angle is specified by the controlling qubits, specifically, the qubits in the supplementary containing α_i (which is shown directly to the right of the downward arrow). This rotation is actually a series of standardized rotations, each controlled by a single qubit of the register storing α_i , as described below.

Due to finite computational resources, the rotation operator must be implemented approximately. Note that since α is a rotation angle between 0 and 2π , so $0 \leq \alpha/2\pi \leq 1$ and we can approximate in k -digits of binary notation that $\alpha/2\pi \approx \sum_{i=1}^k \frac{a_i}{2^i}$ so that $\alpha \approx \sum_{i=1}^k \pi \frac{a_i}{2^{i-1}}$. Note that the a_i are either 1 or 0, as this is a binary representation. The rotation operator is realized as a sequence of at most k standard rotations as $R(\alpha) \approx R(\pi/2^{k-1})^{a_k} \dots R(\pi/2)^{a_2} R(\pi)^{a_1}$. The precision of this operation on the entire array (assuming the standard rotations are exact) is thus $\delta = O(N2^{-k})$ and the size of the circuit required to execute the algorithm is polynomial in $n + \log(1/\delta)$.

PREPARING A GAUSSIAN IN MULTIPLE DIMENSIONS

In S dimensions, the general Gaussian function is $G(\vec{x}) = Ce^{-\vec{x}^T A \vec{x}/2}$, where C is a constant, A is a real symmetric positive definite matrix ($A^T = A$, and $\vec{x}^T A \vec{x} \geq 0 \forall \vec{x} \neq 0$), and \vec{x} is vector with S components, $\vec{x} = \vec{x}' - \vec{\mu}$ where \vec{x}' denotes position in S space and $\vec{\mu}$ is a vector denoting the mean of the distribution.




FIG. 1: Schematic circuit for preparing a Gaussian wavefunction. The top part of the figure shows the general scheme of the recursion. The lower figure is a more detailed view of what occurs in a representative block of the process like the boxed section of the circuit above. The top and bottom multi-qubit wires represent different regions of the semipermanent supplementary memory where all pairs (σ_i, μ_i) are stored until the end of computation. The multi-qubit wires feeding into the black box G are the “working” supplementary memory that is inverted and cleared in each step.

We assume that the register (call it R) contains Sk qubits, where k is some integer. Thus, we can break the array into S subarrays (index them R_i for $1 \leq i \leq S$) each containing k qubits, and write $R=R_1R_2...R_S$, with multiplication representing array concatenation. The state space corresponding to R is of dimension 2^{Sk} . The wavefunction we wish to prepare has form

$$|\psi\rangle = C \sum_{\vec{n} \in B^S} e^{-\vec{n}^T A \vec{n}/2} |\vec{n}\rangle \quad (14)$$

In this expressions $C^2 = \pi^{-S/2} \sqrt{\det A}$ and $\vec{n} \in B^S$, $B = [-2^{k-1}, \dots, 2^{k-1} - 1]$, where negative integers are represented using two's complement. In the quantum computer, $|\vec{n}\rangle$ is the basis vector formed by concatenating all binary representations of the entries of \vec{n} . For example, if $n = 9$, in binary representation, $n = 1001$, so if we assume $k=2$, $\vec{n} = (10, 01)^T = (2, 1)^T$. For R_i , the basis state $|m\rangle$ corresponds to the integer m along the i th coordinate axis in S space. Thus, each subarray generates one coordinate axis of the space in question. Note that this space is not a state space in the system, rather, each integer point in the cube described corresponds to one dimension/basis vector of the system's state space. The reason for the distinction will be clear momentarily.

Our method of preparing the general S -dimensional Gaussian relies on a special case of the problem. Note that if the matrix A in $e^{-\vec{x}^T D \vec{x}/2} = C e^{-\vec{x}^T A \vec{x}/2}$ is diagonal (in which case, we call it D), then $G(\vec{x})$ factors into 1-dimensional Gaussians as $G(\vec{x}) = \prod_{i=1}^N \exp(-d_i x_i^2/2)$ where d_i is the i th diagonal element of D and x_i the i -th entry of \vec{x} .

Thus in this special case, the state to prepare becomes (with $\vec{n}=(n_1, \dots, n_S)$, where the n_i are k -digit binary integers)

$$|\psi\rangle = C \sum_{\vec{n} \in B^S} e^{-\vec{n}^T D \vec{n}/2} |\vec{n}\rangle = C \sum_{\vec{n} \in B^S} \prod_{i=1}^S \exp(-\frac{1}{2} d_i n_i^2) |\vec{n}\rangle = C \bigotimes_{i=1}^S \left(\sum_{n \in B} \exp(-\frac{1}{2} d_i n^2) |n\rangle \right) \quad (15)$$

This state is just a product of 1-dimensional Gaussians, which we have already shown how to prepare. We can prepare this state by preparing N simple Gaussians on k qubit subarrays. The overall normalization constant (here denoted by C) is accounted for by the normalization constants of the Gaussians on subarrays.

Some simple properties of linear algebra and a change of variables allow us to transform this easily prepared state into any arbitrary Gaussian. Any real symmetric positive definite matrix A can be decomposed as $A = (M^T)^{-1} D (M^{-1})$ for some D and M , where D is diagonal and M is upper triangular with 1's on the diagonal (in other words, a multidimensional shearing matrix). Thus, if we apply transformation $\vec{n} \rightarrow \vec{x} = M \vec{n} \Leftrightarrow \vec{n} = M^{-1} \vec{x}$, then

$$\exp(-\vec{n}^T D \vec{n}/2) = \exp\left(-\frac{\vec{x}^T (M^T)^{-1} D M^{-1} \vec{x}}{2}\right) = \exp\left(-\frac{\vec{x}^T A \vec{x}}{2}\right) \quad (16)$$

The matrix M can be decomposed as the product of shearing matrices with exactly one non-zero off-diagonal entry each. Note that M acts numerically on the basis vectors in a classically reversible fashion (provided relevant parameters are stored). It is not a unitary operator and is not applied directly to the state space of the system. Rather, it maps a vector \vec{n} to a vector \vec{x} (and the corresponding basis vector $|\vec{n}\rangle$ to $|\vec{x}\rangle$) as in the following example.

Let M be a square, 2 by 2 upper triangular matrix with 1's the diagonal and a real number α in the upper right entry and let $\vec{n} = (n_1, n_2)^T$. Then, we see that

$$M = \begin{pmatrix} 1 & \alpha \\ 0 & 1 \end{pmatrix}, \quad \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = M \vec{n} = \begin{pmatrix} 1 & \alpha \\ 0 & 1 \end{pmatrix} \begin{pmatrix} n_1 \\ n_2 \end{pmatrix} = \begin{pmatrix} n_1 + \alpha n_2 \\ n_2 \end{pmatrix} \approx \begin{pmatrix} n_1 + \lfloor \alpha n_2 \rfloor \\ n_2 \end{pmatrix} \quad (17)$$

$$\begin{pmatrix} n_1 \\ n_2 \end{pmatrix} = \begin{pmatrix} x_1 - \lfloor \alpha x_2 \rfloor \\ x_2 \end{pmatrix} \quad (18)$$

Thus, M can act on basis vectors numerically by matrix multiplication with rounding, which, provided the specifications for M are stored, is an efficiently reversible classical operation that can thus be carried out on a quantum computer. Applying M to a superposition of states results in a changes of variable of the kind described above, which turns our set of independent Gaussians into the general Gaussian,

$$C \sum_{\vec{n} \in B^S} e^{-\vec{n}^T D \vec{n}/2} |\vec{n}\rangle \rightarrow C \sum_{\vec{x} \in B^S} e^{-\vec{x}^T A \vec{x}/2} |\vec{x}\rangle. \quad (19)$$

Thus, we have an algorithm for preparing the general Gaussian function in N dimensions.

RESAMPLING

One application of the Gaussian states described above is multidimensional resampling. This is more clear after a description of resampling in one dimension. The goal is to stretch or contract a wavefunction with respect to some function f , see eqn. (1). (We assume that f is smooth, monotone, and its derivative is bounded from above and from below.) Such deformations are easily accomplished, provided that ψ is a slowly varying function ($\psi(x)$ and $\psi(x \pm 1)$ are very close for all x in the range of interest). Assume we have two qubit ensembles, A and B , each containing N qubits. The basis states of ensemble A correspond to the x axis, those of ensemble B to the y axis. The desired original wavefunction ψ is prepared on ensemble A , while ensemble B is prepared in state $|0^N\rangle$. The state of the combined system is thus

$$|\eta_0\rangle = \sum_{x=0}^{2^N-1} \psi(x) |x, 0^N\rangle \quad (20)$$

If we want to perform a resampling with respect to $f(x) = \lfloor \frac{x}{a} \rfloor$, then we prepare the uniform superposition state

$$|\kappa_n\rangle = \frac{1}{\sqrt{2n}} \sum_{j=-n}^{n-1} |j\rangle \quad (21)$$

on ensemble B using an algorithm described in Problem 9.4 of [4]. We describe how parameter n is selected shortly. After this operation, the whole ensemble is in state

$$|\eta_1\rangle = \frac{1}{\sqrt{2n}} \sum_{x=0}^{2^N-1} \psi(x) \sum_{j=-n}^{n-1} |x, j\rangle \quad (22)$$

A simple reversible computation can efficiently transform this state into

$$|\eta_2\rangle = \frac{1}{\sqrt{2n}} \sum_{x=0}^{2^N-1} \psi(x) \sum_{j=-n}^{n-1} \left| x, \left\lfloor \frac{x}{a} \right\rfloor + j \right\rangle \quad (23)$$

Recall that we wish to prepare the state $\xi(y) = \sqrt{a}\psi(f(ay))$ on an N -qubit ensemble. Note that if we started with that wavefunction prepared on ensemble B with ensemble A initialized to $|0^N\rangle$, we could carry out a transformation analogous to that above to yield the state

$$|\tilde{\eta}_2\rangle = \frac{1}{\sqrt{2n}} \sum_{y=0}^{2^N-1} \psi(ay) \sum_{j=-an}^{an-1} |[ay] + j, y\rangle. \quad (24)$$

If we think of A and B as generating coordinate axes x and y , then wavefunctions can be visualized as sets of points in the xy plane, specifically, as points occupying the vertices of the infinite grid with integer spacing. In this scheme, $|\eta_2\rangle$ and $|\tilde{\eta}_2\rangle$ appear as bands of width $2n$ about the line $y = \frac{x}{a}$. The amplitudes of basis vectors in $|\eta_2\rangle$ and $|\tilde{\eta}_2\rangle$ agree on that line; the only difference is that the amplitudes in $|\eta_2\rangle$ are constant on vertical strips whereas the amplitudes in $|\tilde{\eta}_2\rangle$ are constant on horizontal strips. If n is much larger than the grid spacing but small enough relative to the grid length at which ψ and ξ vary, the amplitudes in intersecting horizontal and vertical strips are close. Thus we have the approximate equality $|\eta_2\rangle \cong |\tilde{\eta}_2\rangle$. This means that by choosing N and n appropriately and preparing the state $|\eta_2\rangle$, it is virtually equivalent to having prepared the state $|\tilde{\eta}_2\rangle$. If we then apply (to the combined ensemble) the inverse of the computation we would have used to prepare $|\tilde{\eta}_2\rangle$, it is as if we prepared and then uncomputed $|\tilde{\eta}_2\rangle$, leaving the combined ensemble in (or rather, arbitrarily close to) the state such that the wavefunction on B is $\xi(y)$ and the wavefunction on A is $|0^N\rangle$. Thus if we can prepare $\psi(x)$ satisfying the necessary smoothness conditions, we can use resampling with supplementary qubits to transform the state to one arbitrarily close to $\xi(y)$.

This particular transformation takes advantage of the fact that uniform superpositions on an interval are invariant under stretching or contracting transformation. Multidimensional resampling requires a family of states that is similarly invariant under a general linear transformation. Multidimensional Gaussian states satisfy this condition and thus may be used for multidimensional resampling.




FIG. 2: Conceptual representation of states $|\tilde{\eta}_2\rangle$ and $|\eta_2\rangle$. The grey region (and the dotted region in the inset) represents the basis states of the combined ensemble AB that have non-zero amplitudes. In state $|\eta_2\rangle$ the amplitudes on vertical strips are equal, in state $|\tilde{\eta}_2\rangle$ amplitudes are constant on horizontal strips. If n is much larger than the spacing between basis vectors but much smaller than the size of the grid, the amplitudes of basis vectors in a vertical strip will be roughly equal to those of the basis vectors of any horizontal strip that intersects it. When this is the case, $|\tilde{\eta}_2\rangle \cong |\eta_2\rangle$ and resampling will work.

We illustrate with a simple example. Note that in the procedure above, if we prepare on ensemble B the state $|\xi_{\sigma,\mu,N}\rangle$ with parameters σ and μ such that

$$|\xi_{\sigma,\mu,N}\rangle \cong \sum_{i=0}^{2^N-1} e^{-\frac{(i-\mu)^2}{2\sigma^2}} |i\rangle \cong |\xi_{\sigma,\mu}\rangle = \sum_{i=-\infty}^{\infty} e^{-\frac{(i-\mu)^2}{2\sigma^2}} |i\rangle \quad (25)$$

with $\mu \gg \sigma \gg 1$ and $\sigma \cong 2n$ instead of the uniform superposition state $|\kappa_n\rangle$, the procedure still works. Intuitively, we prepare a series of vertical, parallel Gaussian-weighted strips centered somewhere above the x axis and shift them up along the line $y = \frac{x}{a}$ so that each strip is centered at the line. The result is that the entire grid has non-zero amplitudes, but all amplitudes except those immediately about $y = \frac{x}{a}$ are very, very small, so the result looks very similar to Figure 2 and is amenable to resampling as described.

ACKNOWLEDGEMENTS

This research was funded by Caltech's Summer Undergraduate Research Fellowship (SURF) Program. W.W. would like to thank Ning Bao, Erik Winfree, and David Renshaw for helpful comments and discussion.

-
- [1] P. W. Shor. Algorithms for Quantum Computation: Discrete Logarithms and Factoring. IN Proceedings, 35th Annual Symposium on Foundations of Computer Science, IEEE Press, Los Alamitos, CA, 1994.
 - [2] L. K. Grover. Quantum Mechanics Helps in Searching for a Needle in a Haystack. Phys. Rev. Lett., 79(2):325, 1997.
 - [3] K. Ireland, M. Rosen. A Classical Introduction to Modern Number Theory. Graduate Texts in Mathematics, v. 84. Springer, 1998 (2nd edition).

- [4] A. Yu Kitaev, A. H. Shen, M. N. Vyalyi. Classical and Quantum Computation. Graduate Studies in Mathematics, v. 47. American Mathematical Society, 2002.
- [5] L. K. Grover, T. Rudolph. Creating superpositions that correspond to efficiently integrable probability distributions. quant-ph/0208112.